

Greedy Feature Selection for Subspace Clustering

Eva L. Dyer

E.DYER@RICE.EDU

*Department of Electrical & Computer Engineering
Rice University, Houston, TX, 77005, USA*

Aswin C. Sankaranarayanan

SASWIN@ECE.CMU.EDU

*Department of Electrical & Computer Engineering
Carnegie Mellon University, Pittsburgh, PA, 15213, USA*

Richard G. Baraniuk

RICHB@RICE.EDU

*Department of Electrical & Computer Engineering
Rice University, Houston, TX, 77005, USA*

Editor:

Abstract

Unions of subspaces are powerful nonlinear signal models for collections of high-dimensional data. However, existing methods that exploit this structure require that the subspaces the signals of interest occupy be known a priori or be learned from the data directly. In this work, we analyze the performance of greedy feature selection strategies for learning unions of subspaces from an ensemble of data. We develop sufficient conditions that are required for orthogonal matching pursuit (OMP) to recover subsets of points from the ensemble that live in the same subspace, a property we refer to as *exact feature selection* (EFS). Following this analysis, we provide an empirical study of greedy feature selection strategies for subspace clustering and characterize the gap between sparse recovery methods and nearest neighbor (NN)-based approaches. We demonstrate that the gap between sparse recovery and NN methods is particularly pronounced when the tiling of subspaces in the ensemble is sparse, suggesting that sparse recovery methods can be used in a number of regimes where nearest neighbor approaches fail to reveal the subspace membership of points in the ensemble.

Keywords: Subspace clustering, unions of subspaces, hybrid linear models, sparse approximation, structured sparsity, nearest neighbors, low-rank approximation.

1. Introduction

1.1 Unions of Subspaces

With the emergence of novel sensing systems capable of acquiring data at scales ranging from the nano to the peta, modern sensor and imaging data are becoming increasingly high-dimensional and heterogeneous. To cope with this explosion of high-dimensional data, one must exploit the fact that low-dimensional geometric structure exists amongst collections of data.

Linear subspace models are one of the most widely used signal models for collections of high-dimensional data, with applications throughout signal processing, machine learning, and the computational sciences. This is due in part to the simplicity of linear models but also due to the fact that principal components analysis (PCA) provides a closed-form and

computationally efficient solution to the problem of finding an optimal low-rank approximation to a collection of data (an ensemble of points in \mathbb{R}^n). More formally, if we stack a collection of d vectors (points) in \mathbb{R}^n into the columns of $Y \in \mathbb{R}^{n \times d}$, then PCA finds the best rank- k estimate of Y by solving

$$\text{(PCA)} \quad \min_{X \in \mathbb{R}^{n \times d}} \|Y - X\|_F \quad \text{subject to} \quad \text{rank}(X) \leq k. \quad (1)$$

In many cases, a linear subspace model is sufficient to characterize the intrinsic structure of the ensemble; however, in many emerging applications, a single subspace is not enough. Instead, ensembles must be modeled as living on a *union of subspaces* or a union of affine planes of mixed or equal dimension. Ensembles ranging from collections of images taken of objects under different illumination conditions (Basri and Jacobs, 2003; Ramamoorthi, 2002), motion trajectories of point-correspondences (Kanatani, 2001), to structured sparse and block-sparse signals (Lu and Do, 2008; Blumensath and Davies, 2009; Baraniuk et al., 2010) are all well-approximated by a union of low-dimensional subspaces or a union of affine hyperplanes. Union of subspace models have also found utility in the classification of signals collected from complex and adaptive systems at different instances in time, e.g., electrical signals collected from the brain’s motor cortex (Gowreesunker et al., 2011).

Unions of subspaces provide a natural extension to single subspace models, but providing an extension of PCA that leads to *provable* guarantees for learning multiple subspaces is challenging. This is due to the fact that segmentation—the identification of points that live in the same subspace—and subspace estimation must be performed simultaneously. However, if we can accurately sift through the points in the ensemble and determine which points lie along or near the same subspace, then subspace estimation becomes trivial. For this reason, many state-of-the art methods for learning unions of subspaces rely on first forming *local subspace estimates*¹ from a subset of points in the data.

A common heuristic used to obtain local subspace estimates is to select points that lie within an Euclidean neighborhood of one another (or a fixed number of nearest neighbors (NN)) and then form a local estimates from multiple sets of NNs. At a high-level, most NN-based approaches for subspace clustering can be summarized as consisting roughly of the following three steps:

- (1) For the i^{th} point in the set, y_i , select a set of points from the ensemble that live within an ϵ -radius from y_i in terms of their Euclidean distance. Denote this subset of points Y_Λ , where Λ is an index set containing the indices of all the neighbors of y_i .
- (2) Form a low-rank PCA estimate by solving (1) for the points in the sub-matrix Y_Λ .
- (3) Compute the *subspace affinity matrix* $W \in \mathbb{R}^{d \times d}$ for the ensemble, where the (i, j) entry of the matrix represents the likelihood that y_i and y_j live close to the same subspace or whether y_i and y_j produce similar local subspace estimates.

Methods that use NN sets to form local subspace estimates from the data include local subspace affinity (LSA) (Yan and Pollefeys, 2006), spectral clustering based on locally

1. A local subspace estimate is a low-rank approximation formed from a subset of points in the ensemble, rather than from the entire collection of data.

linear approximations (Arias-Castro et al., 2010), spectral curvature clustering (Chen and Lerman, 2009), and local best-fit flats (Zhang et al., 2010). The main differences between these methods lie either in the way that the entries of the affinity matrix are computed in Step 3 or the way in which this matrix is used to obtain an estimate of the underlying subspace structures present in the ensemble. In the case of approaches built upon spectral clustering (Shi and Malik, 2000), one performs spectral clustering on the subspace affinity matrix for the ensemble in order to cluster the data into different subspaces. In the case of consensus-based approaches, one finds a robust estimate of the mode of the local subspace estimates formed in Step 2; this problem can also be posed as an optimization on the subspace affinity matrix.

When the subspaces present in the ensemble are linearly separable or non-intersecting, local subspace estimates formed from NNs provide relatively reliable and stable estimates of the subspaces present in the ensemble. However, neighborhood-based approaches quickly fail as the separation between the two structures decreases and as the subspace dimension increases relative to the number of points in each subspace. This is due in part to the fact that, as the dimension of the intersection between two subspaces increases, the Euclidean distance between points becomes a poor predictor of which points belong to the same subspace. Thus, we seek an alternative strategy for forming local estimate that does not rely solely on whether points in the same subspace live in a local Euclidean neighborhood. Instead, our goal is to identify another strategy for “feature selection” that returns sets of points (feature sets) that lie along the same subspace.

1.2 Exact Feature Selection

Instead of computing local subspace estimates from sets of NNs, Elhamifar and Vidal (2009) propose a novel approach for feature selection based upon forming sparse representations of the data via ℓ_1 -minimization. The main intuition underlying their approach is that when a sparse representation of a point contained within a union of subspaces is formed with respect to the remaining points in the dataset, the representation should only consist of other points that belong to the same subspace. When a sparse representation consists of points that lie in the same subspace, we say that *exact feature selection* (EFS) occurs. Under certain assumptions on both the sampling and “distance between subspaces”,² this approach to feature selection leads to provable guarantees that EFS will occur (Elhamifar and Vidal, 2010; Soltanolkotabi and Candès, 2012), even when the subspaces intersect.

We refer to this application of sparse recovery as *endogenous sparse recovery* due to the fact that representations are not formed from an external collection of primitives (such as a basis or dictionary) but are formed “from within” the data. Formally, for a set of d signals $\mathcal{Y} = \{y_1, \dots, y_d\}$, each of dimension n , the sparsest representation of the i^{th} point y_i is defined as

$$c_i^* = \arg \min_{c \in \mathbb{R}^d} \|c\|_0 \quad \text{subject to} \quad y_i = \sum_{j \neq i} c(j)y_j, \quad (2)$$

where $\|c\|_0$ counts the number of non-zeroes in its argument. Let $\Lambda^{(i)} = \text{supp}(c_i^*)$ denote the subset of points selected to represent the i^{th} point and $c_i^*(j)$ denote the contribution

2. The distance between a pair of subspaces is typically measured with respect to the principal angles between the subspaces or other related distances on the Grassmanian manifold.

of the j^{th} point to the endogenous representation of y_i . By penalizing representations that require a large number of non-zero coefficients, the resulting representation will be sparse.

In general, finding the sparsest representation of a signal has combinatorial complexity, and so sparse recovery methods such as basis pursuit (BP) (Chen et al., 1998) or a greedy pursuit such as orthogonal matching pursuit (OMP) (Davis et al., 1994) are employed to find approximate solutions.

1.3 Contributions

In Elhamifar and Vidal (2010), the authors show that when subspaces are *disjoint* (intersect only at the origin) and the minimum principal angle between subspaces is sufficiently large, the points selected by BP will belong to the same subspace, i.e., EFS is guaranteed. Recently, Soltanolkotabi and Candès (2012) developed both deterministic and probabilistic guarantees for EFS from unions of intersecting subspaces.

In parallel with recent developments for subspace clustering with BP, in this paper, we provide an analysis of endogenous sparse recovery with OMP. The main result of our analysis is a new geometric condition (Thm. 1) for EFS that highlights the tradeoff between the *mutual coherence* between points living in different subspaces and the *covering radius* of the points within a common subspace. The covering radius can be interpreted as the radius of the largest ball that can be embedded within each normalized subspace cluster without touching a point in the ensemble; the vector that lies at the center of this open ball, or the vector in the subspace that attains the covering radius is referred to as a *deep hole*. Thm. 1 suggests that subspaces can be arbitrarily close to one another and even intersect, as long as the data is distributed “nicely” along each subspace. By “nicely”, we mean that the points that tile each subspace, attain a small covering radius and do not cluster together, leaving large gaps in the subspace. In Fig. 1, we illustrate the covering radius of a set of points on the sphere (the deep hole is denoted by a star).

After introducing a general geometric condition for EFS with OMP, we extend this analysis to the case where the data live on what we refer to as an *uniformly bounded union* of subspaces. This analysis demonstrates that when the points living in a particular subspace are incoherent with the principal vectors that support pairs of subspaces in the ensemble, EFS can be guaranteed; this is true even when subspaces in the ensemble are intersecting. Our condition for bounded subspaces suggests that, in addition to properties related to the sampling of the subspace, one can characterize the separability of pairs of subspaces by examining the correlation between the data and the unique set of principal vectors that support pairs of subspaces in the ensemble. The principal vectors have a straightforward interpretation: the principal vectors are the collection of vectors from a pair of subspaces that are maximally correlated with one another but also form an orthonormal basis (ONB) for each of their respective subspaces.

In addition to providing a theoretical analysis of EFS, another major contribution of this work is revealing the “gap” between nearest neighborhood-based (NN) approaches and sparse recovery methods such as OMP and BP for feature selection. In particular, we demonstrate that sparse recovery methods admit a higher proportion of points that belong to the same subspace than sets of NN; this gap in performance is particularly pronounced when the number of points in each subspace decreases. These empirical results point to




Figure 1: *Covering radius of points in a normalized subspace.* The interior of the antipodal convex hull of points in a normalized subspace—a subspace of \mathbb{R}^n mapped to the unit ℓ_2 -sphere—is shaded. The vector in the normalized subspace (circle) that attains the covering radius or lies within a deep hole in the subspace is denoted with a star: when compared with the convex hull, the deep hole coincides with the maximal gap between the convex hull and the set of all vectors that live in the normalized subspace.

another advantage of forming sparse representations from within the data; sparse recovery methods provide a natural way to reveal the subspace affinity amongst points that might be far away from one another in a Euclidean sense. By exploiting non-local relationships between points, sparse recovery methods are capable of providing reliable subspace estimates with far fewer points than neighborhood-based estimates.

To illustrate the role that the sampling density has on subspace clustering in real-world datasets, we compare the performance of NN and sparse recovery methods (BP and OMP) for clustering unlabeled faces that have been captured under different illumination conditions—the separation of points that live on different “illumination subspaces”. We provide empirical evidence that while both NN and sparse recovery methods have comparable rates of EFS for dense samplings of each subspace, when the subspaces are more sparsely sampled, sparse recovery methods provide significant advantages over NN, both in terms of EFS and in their final classification rates. Moreover, we demonstrate that in a number of settings, OMP outperforms BP in terms of both EFS and clustering performance. See Fig. 2 for an example of the affinity matrices formed from pairs of face subspaces for NN (left), BP (middle), and OMP (right).




Figure 2: *Comparison of subspace affinity matrices for illumination subspaces.* In each row, we display the subspace affinity matrices obtained for a different pair of illumination subspaces in the dataset, for NN (left), BP (middle), and OMP (right). To the left of the affinity matrices, we display an exemplar image from each illumination subspace.

1.4 Paper Organization

We now provide a roadmap for the rest of the paper.

Section 2. We introduce our signal model, the sparse subspace clustering (SSC) algorithm introduced in (Elhamifar and Vidal, 2009), and describe how OMP may be used for feature selection in subspace clustering.

Section 3 and 4. We develop the main theoretical results of this paper and provide new geometric insights into EFS from unions of subspaces. We introduce sufficient conditions for EFS to occur with OMP for general unions of subspaces in Thm. 1, disjoint unions in Cor. 1, and uniformly bounded unions in Thm. 3.

Section 5. We conduct a number of numerical experiments to validate our theory and compare feature selection with sparse recovery methods to NN-based feature selection. Experiments are provided for both synthetic and real data.

Section 6. We discuss the implications of our theoretical analysis and subsequent studies on sparse approximation, dictionary learning, and compressive sensing. We conclude with a number of interesting open questions and future lines of research.

Section 7. We supply the proofs of the results contained in Sections 3 and 4.

1.5 Notation

In this paper, we will work solely in real finite-dimensional vector spaces, \mathbb{R}^n . We write vectors x in lowercase script, matrices A in uppercase script, and scalar entries of vectors as $x(j)$. The standard p -norm is defined as

$$\|x\|_p = \left(\sum_{j=1}^n |x(j)|^p \right)^{1/p},$$

where $p \geq 1$. The ℓ_0 quasi-norm of a vector x is defined as the number of non-zero elements in x . The support of a vector x , often written as $\text{supp}(x)$, is the set containing the indices of its non-zero coefficients; hence, $\|x\|_0 = |\text{supp}(x)|$. We denote the Moore-Penrose pseudoinverse of a matrix A as A^\dagger . If $A = U\Sigma V^T$ then $A^\dagger = V\Sigma^+U^T$, where we obtain Σ^+ by taking the reciprocal of the entries in Σ , leaving the zeros in their places, and taking the transpose. An orthonormal basis (ONB) Φ that spans the subspace \mathcal{S} of dimension k satisfies the following two properties: $\Phi^T\Phi = I_k$ and $\text{range}(\Phi) = \mathcal{S}$, where I_k is the $k \times k$ identity matrix. Let $P_\Lambda = X_\Lambda X_\Lambda^\dagger$ denote an ortho-projector onto the subspace spanned by the sub-matrix X_Λ .

2. Greedy Feature Selection for Subspace Clustering

In this section, we introduce our signal model, detail the sparse subspace clustering (SSC) method developed in Elhamifar and Vidal (2009), and discuss the utility of greedy methods for feature selection in subspace clustering.

2.1 Signal Model

Given a set of p subspaces of \mathbb{R}^n , each of dimension $k_i \leq k$, which we denote as $\mathcal{S}_1, \dots, \mathcal{S}_p$, we generate “subspace clusters” by sampling d_i points from each subspace \mathcal{S}_i . Let $\tilde{\mathcal{Y}}_i$ denote the resulting set of points from subspace \mathcal{S}_i and let $\tilde{\mathcal{Y}} = \cup_{i=1}^p \tilde{\mathcal{Y}}_i$ denote the union of these p subspace clusters.

To normalize the points in each subspace cluster or generate a “normalized subspace cluster”, we map each point in $\tilde{\mathcal{Y}}$ to the unit sphere. Let

$$\mathcal{Y} = \left\{ \frac{y_1}{\|y_1\|_2}, \frac{y_2}{\|y_2\|_2}, \dots, \frac{y_d}{\|y_d\|_2} \right\}$$

denote the resulting set of unit norm points and \mathcal{Y}_i denote the set of unit norm points that lie in the span of subspace \mathcal{S}_i . Let $\mathcal{Y}_{-i} = \mathcal{Y} \setminus \mathcal{Y}_i$ denote the set of points in \mathcal{Y} , where the points in \mathcal{Y}_i are excluded.

Let $Y = [Y_1 \ Y_2 \ \cdots \ Y_p]$ denote the matrix of normalized data, where each point in \mathcal{Y}_i is stacked into the columns of $Y_i \in \mathbb{R}^{n \times d_i}$. The points in Y_i can be expanded in terms of an orthonormal basis $\Phi_i \in \mathbb{R}^{n \times k_i}$ that spans \mathcal{S}_i and subspace coefficients $A_i = \Phi_i^T Y_i$, where $Y_i = \Phi_i A_i$. Let Y_{-i} denote the matrix containing the points in Y with the submatrix Y_i excluded.

2.2 Sparse Subspace Clustering

In (Elhamifar and Vidal, 2009), the authors propose a novel approach to feature selection called sparse subspace clustering (SSC) that employs a relaxation of the ℓ_0 -minimization problem in (2). To be precise, the sparse subspace clustering (SSC) algorithm proceeds by solving the following problem for each point in \mathcal{Y} :

$$c_i^* = \arg \min_{c \in \mathbb{R}^d} \|c\|_1 \quad \text{subject to} \quad y_i = \sum_{j \neq i} c(j) y_j. \quad (3)$$

After finding the solution to this ℓ_1 -minimization problem, each d -dimensional feature vector c_i^* is placed into the i^{th} row or column of the affinity matrix C and spectral clustering (Shi and Malik, 2000) is performed on the graph Laplacian of $W = |C| + |C^T|$.

In (Elhamifar and Vidal, 2012), the authors provide an extension of SSC to the case where the data might not admit an exact representation with respect to other points in the ensemble. In this case, they employ an inequality constrained version of BP known as basis pursuit denoising (BPDN) for feature selection; for each point y_i , they solve the following problem

$$c_i^* = \arg \min_{c \in \mathbb{R}^d} \|c\|_1 \quad \text{subject to} \quad \|y_i - \sum_{j \neq i} c(j) y_j\|_2 < \kappa, \quad (4)$$

where κ is a parameter that is selected based upon the amount of noise in the data. A robust variant of this algorithm was recently proposed in Soltanolkotabi et al. (2013) to perform sparse subspace clustering on data living on noisy subspaces.

2.3 Greedy Feature Selection

Instead of solving the endogenous sparse recovery problem in (2) via ℓ_1 -minimization, as originally proposed in SSC, we will study the behavior of greedy pursuit (OMP) for endogenous recovery. We detail this procedure in Alg. 1.

For each point y_i , the output of the OMP algorithm is a *feature set*, $\Lambda^{(i)}$, which indexes the columns in Y that are used to form a representation of y_i . To use these feature sets to cluster the data as in SSC, a d -dimensional sparse feature vector \bar{c}_i is computed by stacking the k -dimensional projection $c_i = Y_{\Lambda^{(i)}}^\dagger y_i$ into the entries of \bar{c}_i indexed by $\Lambda^{(i)}$, where $Y_{\Lambda^{(i)}}^\dagger \in \mathbb{R}^{k \times n}$ is the pseudoinverse of the submatrix $Y_{\Lambda^{(i)}} \in \mathbb{R}^{n \times k}$. The feature vector \bar{c}_i is then stacked in the i^{th} row of $C \in \mathbb{R}^{d \times d}$ and the subspace affinity matrix of the ensemble is computed as $W = |C| + |C^T|$.

After computing the subspace affinity matrix W with OMP, spectral clustering may be performed on the graph Laplacian of W (as in SSC), or a consensus-based method may be

Algorithm 1 : Greedy Feature Selection with OMP

Input: Input signal y_i , a matrix containing d data points $Y \in \mathbb{R}^{n \times d}$, and a stopping criterion.

Output: A feature set $\Lambda^{(i)}$ containing the indices of all points chosen in the pursuit.

Initialize: Set the residual to the input signal $r = y_i$.

1. Select the point that is maximally correlated with the residual and add it to $\Lambda^{(i)}$

$$\Lambda^{(i)} \leftarrow \Lambda^{(i)} \cup \arg \max_{j \neq i} |\langle y_j, r \rangle|.$$

2. Update the residual by projecting the signal into the space orthogonal to the span of the points in $\Lambda^{(i)}$

$$r \leftarrow (I - Y_{\Lambda^{(i)}} Y_{\Lambda^{(i)}}^\dagger) y_i.$$

3. Repeat steps (1)–(2) until the stopping criterion is reached, e.g., either $|\Lambda^{(i)}| = k$ or the norm of the residual $\|r\| \leq \delta$.
-

employed. In (Dyer, 2011), we introduce a consensus-based algorithm for subspace clustering and provide a comparison of consensus and clustering-based approaches for subspace clustering with greedy feature selection.

Although OMP is known to be suboptimal for standard applications of sparse signal recovery, our results suggest that greedy pursuits provide a powerful alternative to convex optimization based approaches such as BP and BPDN. An obvious advantage of using greedy methods is that they exhibit reduced computational complexity when compared to convex optimization-based approaches. In addition, we observe that OMP recovers representations that “fill in” subspace affinity matrices more uniformly than BP. This in turn results in affinity matrices that are easier to segment—as measured by the ratio of the sum of the edges within a cluster to the sum of the edges across different clusters. See Fig. 2 for an example of the affinity matrices obtained via OMP, BP, and NN for collections of images of faces under different lighting conditions in the Yale B database (Georghiades et al., 2001).

In Section 5.4, we provide empirical evidence that sparse recovery methods (i.e., OMP and BP) provide significant improvements over NN in terms of EFS and the clustering performance of the three methods on real and synthetic data. These empirical findings, coupled with the fact that OMP exhibits reduced computational complexity, suggest that OMP provide significant advantages over BP in spectral clustering-based approaches to subspace clustering.

3. Exact Feature Selection from Unions of Subspaces

In this section, we provide a formal definition of EFS and develop sufficient conditions that guarantee that EFS will occur for all of the points contained within a particular subspace cluster.

3.1 Exact Feature Selection

As we alluded to in the Introduction, in order to guarantee that OMP returns a sample set that yields accurate local subspace estimates, we will be interested in determining when the feature set $\Lambda^{(i)}$ returned by Alg. 1 only contains other points that belong to the same subspace. When the feature set only contains points that live in the same subspace, we say that *exact feature selection* (EFS) occurs. We now supply a formal definition.

Definition 1 (Exact feature selection) *Let $\mathcal{Y}_k = \{y : (I - P_k)y = 0, y \in \mathcal{Y}\}$ index the set of points in \mathcal{Y} that live in the span of subspace \mathcal{S}_k , where P_k is a projector onto the span of subspace \mathcal{S}_k . For a point $y_i \in \mathcal{Y}_k$ with feature set $\Lambda^{(i)}$, we say that $\Lambda^{(i)}$ contains exact features if $y_j \in \mathcal{Y}_k, \forall j \in \Lambda^{(i)}$.*

EFS is essential for studying the performance of algorithms for unsupervised subspace learning problems, because when EFS occurs for a point in the set, this will yield a subspace estimate that coincides with one of the true subspaces contained within the data. For this reason, EFS provides a natural condition for studying the performance of both subspace consensus and spectral clustering methods. Note that this definition allows for points that lie in the intersection between two planes to belong to either subspace. Thus, EFS can still be satisfied for points at an intersection granted they admit feature sets that are restricted to a single subspace cluster.

3.2 Geometric Conditions for EFS

We will now develop our main results which provide sufficient conditions for EFS to occur for all points in a particular subspace cluster.

3.2.1 PRELIMINARIES

Our main result in Thm. 1 below requires measures of both the distance between points in *different subspace clusters* and in the *same subspace cluster*. A natural measure of the similarity between points living in different subspaces is the *mutual coherence*. A formal definition of the mutual coherence is provided below in Def. 2.

Definition 2 *The mutual coherence between the points in the sets $(\mathcal{Y}_i, \mathcal{Y}_j)$ is defined as*

$$\mu_c(\mathcal{Y}_i, \mathcal{Y}_j) = \max_{u \in \mathcal{Y}_i, v \in \mathcal{Y}_j} |\langle u, v \rangle|, \text{ where } \|u\|_2 = \|v\|_2 = 1. \quad (5)$$

In words, the mutual coherence provides a point-wise measure of the normalized inner product (coherence) between all pairs of points that lie in two different subspace clusters.

Let $\mu_c(\mathcal{Y}_i)$ denote the maximum mutual coherence between the points in \mathcal{Y}_i and all other subspace clusters in the ensemble, where

$$\mu_c(\mathcal{Y}_i) = \max_{i \neq j} \mu_c(\mathcal{Y}_i, \mathcal{Y}_j).$$

A related quantity that provides an upper bound on the mutual coherence between points in two subspaces is the cosine of the first *principal angle* between the subspaces. The first principal angle θ_{ij}^* between subspaces \mathcal{S}_i and \mathcal{S}_j of dimension k_i and k_j respectively, is

the smallest angle between a pair of unit vectors (u_1, v_1) drawn from $\mathcal{S}_i \times \mathcal{S}_j$. Formally, the first principal angle is defined as

$$\theta_{ij}^* = \min_{u \in \mathcal{S}_i, v \in \mathcal{S}_j} \arccos \langle u, v \rangle \quad \text{subject to} \quad \|u\|_2 = 1, \|v\|_2 = 1. \quad (6)$$

Whereas the mutual coherence provides a measure of the pairwise similarity between a *finite set* of unit norm vectors that live in different subspaces, the cosine of the minimum principal angle provides a measure of the pairwise similarity between *any pair* of unit norm vectors that lie in the span of $\mathcal{S}_i \times \mathcal{S}_j$. For this reason, the cosine of the first principal angle provides an upper bound on the mutual coherence. The following upper bound is in effect for each pair of subspace clusters in the ensemble:

$$\mu_c(\mathcal{Y}_i, \mathcal{Y}_j) \leq \cos(\theta_{ij}^*). \quad (7)$$

To measure the degree to which the points in the same subspace cluster provide a tiling of their span, we will study the covering radius of each subspace cluster relative to the projective distance. Formally, the covering radius of the set \mathcal{Y}_k is defined as

$$\text{cover}(\mathcal{Y}_k) = \max_{u \in \mathcal{S}_k} \min_{y \in \mathcal{Y}_k} \text{dist}(u, y), \quad (8)$$

where the projective distance between two vectors u and y is defined relative to the acute angle between the vectors

$$\text{dist}(u, y) = \sqrt{1 - \frac{|\langle u, y \rangle|^2}{\|u\|_2^2 \|y\|_2^2}}. \quad (9)$$

The covering radius of the normalized subspace cluster \mathcal{Y}_i can be interpreted as the size of the largest open ball that can be placed in the set of all unit norm vectors that lie in the span of \mathcal{S}_i , without touching a point in \mathcal{Y}_i .

Let (u_i^*, y_i^*) denote a pair of vectors that attain the maximum covering diameter for \mathcal{Y}_i , where $u_i^* \in \mathcal{S}_i$ lies in a deep hole in \mathcal{Y}_i along \mathcal{S}_i . The covering radius can be interpreted as the sine of the angle between the deep hole u_i^* and its nearest neighbor $y_i^* \in \mathcal{Y}_i$. We show the geometry underlying the covering radius in Figure 1.

In the sequel, we will be interested in the maximum (worst-case) covering attained over all d_i sets formed by removing a single point from \mathcal{Y}_i . We supply a formal definition in Def. 3.

Definition 3 *The maximum covering diameter ϵ of the set \mathcal{Y}_i along the subspace \mathcal{S}_i is defined as*

$$\epsilon = \max_{j=1, \dots, d_i} 2 \text{ cover}(\{\mathcal{Y}_i \setminus y_j\}).$$

Hence, the covering radius equals $\epsilon/2$.

A related quantity is the *inradius* of the set \mathcal{Y}_i , or the cosine of the angle between a point in \mathcal{Y}_i and any point in \mathcal{S}_i that attain the covering radius. The relationship between the covering diameter ϵ and inradius $r(\mathcal{Y}_i)$ is given by

$$r(\mathcal{Y}_i) = \sqrt{1 - \frac{\epsilon^2}{4}}. \quad (10)$$

A geometric interpretation of the inradius is that it measures the distance from the origin to the maximal gap in the antipodal convex hull of the points in \mathcal{Y}_i . The geometry underlying the covering radius and the inradius is displayed in Figure 1.

3.2.2 GENERAL RESULT FOR EFS

We are now equipped to state our main geometric result for EFS with OMP. The proof is contained in Section 7.1.

Theorem 1 *Let ϵ denote the maximal covering diameter of the subspace cluster \mathcal{Y}_i as defined in Def. 3. A sufficient condition for EFS to occur for all points in \mathcal{Y}_i is that the mutual coherence*

$$\mu_c(\mathcal{Y}_i) < \sqrt{1 - \frac{\epsilon^2}{4}} - \frac{\epsilon}{\sqrt[4]{12}} \max_{j \neq i} \cos(\theta_{ij}^*), \quad (11)$$

where θ_{ij}^* is the minimum principal angle defined in (6).

In words, this condition requires that the mutual coherence between points in *different subspaces* is less than the difference of two terms that both depend on the covering radius of points along a *single subspace*. The first term on the RHS of (11) is the inradius, as defined in (10); the inradius provides a measure of the coherence between two points that attain the covering radius of the subspace cluster. The second term on the RHS of (11) is the product of the cosine of the minimum principal angle between pairs of subspaces in the ensemble and the covering diameter ϵ of the points in \mathcal{Y}_i .

When subspaces in the ensemble intersect, i.e., $\cos(\theta_{ij}^*) = 1$, condition (11) in Thm. 1 can be simplified to

$$\mu_c(\mathcal{Y}_i) < \sqrt{1 - \frac{\epsilon^2}{4}} - \frac{\epsilon}{\sqrt[4]{12}} \approx \sqrt{1 - \frac{\epsilon^2}{4}} - \frac{\epsilon}{1.86}. \quad (12)$$

In this case, EFS can be guaranteed as long as the points in different subspace clusters are bounded away from intersections between subspaces. When the covering radius shrinks to zero (full sampling of subspace), Thm. 1 requires that $\mu_c < 1$, or that points from different subspaces do not lie exactly in the subspace intersection, i.e., are identifiable from one another.

3.2.3 GEOMETRY UNDERLYING EFS

The main idea underlying the proof of Thm. 1 is that, at each iteration of Alg. 1, we require that the residual used to select a point to be included in the representation is closer to a point in the *correct subspace cluster* (\mathcal{Y}_i) than a point in an *incorrect subspace cluster* (\mathcal{Y}_{-i}). To be precise, we require that the normalized inner product of the residual r and all points outside of the correct subspace cluster

$$\max_{y \in \mathcal{Y}_{-i}} \frac{|\langle r, y \rangle|}{\|r\|_2} < r(\mathcal{Y}_i), \quad (13)$$

at each iteration of Alg. 1. To provide the result in Thm. 1, we require that (13) holds for all possible residuals, or for all $r \in \mathcal{S}_i$.




Figure 3: *Geometry underlying EFS.* A union of two disjoint subspaces of different dimension: the antipodal convex hull of a set of points (red circles) living on a 2D subspace is shaded (green). In (a), we show an example where EFS is guaranteed—the projection of points along the 1D subspace lie inside the shaded convex hull of points in the plane. In (b), we show an example where EFS is not guaranteed—the projection of points along the 1D subspace lie outside the shaded convex hull.

A geometric interpretation of the EFS condition in Thm. 1 is that the projection of points in an incorrect subspace cluster onto the subspace spanned by the correct subspace cluster must lie within the antipodal convex hull of the points in the correct subspace cluster. To see this, consider the projection of the points in incorrect subspace clusters onto subspace \mathcal{S}_i . Let z_j^* denote the point on subspace \mathcal{S}_i that is closest to the signal $y_j \in \mathcal{Y}_{-i}$,

$$z_j^* = \arg \min_{z \in \mathcal{S}_i} \|z - y_j\|_2.$$

We can also write this projection in terms of an orthogonal projection operator $P_i = \Phi_i \Phi_i^T$, where Φ_i is an ONB that spans \mathcal{S}_i and $z_j^* = P_i y_j$.

By definition, the normalized inner product of the residual with points in incorrect subspace clusters is upper bounded as

$$\max_{y_j \in \mathcal{Y}_{-i}} \frac{|\langle r, y_j \rangle|}{\|r\|_2} \leq \max_{y_j \in \mathcal{Y}_{-i}} \frac{|\langle z_j^*, y_j \rangle|}{\|z_j^*\|_2} = \max_{y_j \in \mathcal{Y}_{-i}} \cos \angle \{z_j^*, y_j\} \quad (14)$$

Thus to guarantee EFS, we require that the cosine of the angle between all signals in \mathcal{Y}_{-i} and their projection onto \mathcal{S}_i is less than the inradius of \mathcal{Y}_i . Said another way, the EFS condition requires that the length of all projected points be less than the inradius of \mathcal{Y}_i .

In Fig. 3, we provide a geometric visualization of the EFS condition for a union of disjoint subspaces (union of a 1D subspace with a 2D subspace). In (a), we show a case

where EFS is guaranteed because the projection of points in the incorrect subspace (line) onto the correct subspace (plane) lies well within the antipodal convex hull of the points in the plane. In (b), we show a case where EFS is not guaranteed because the projection of points in the incorrect subspace (line) onto the correct subspace (plane) lies outside the antipodal convex hull of the points in the plane.

3.2.4 EFS FOR DISJOINT SUBSPACES

When the subspaces in the ensemble are *disjoint*, i.e., $\cos(\theta_{ij}^*) < 1$, Thm. 1 can be simplified further by using the bound for the mutual coherence in (7). This simplification results in the following corollary.

Corollary 1 *Let θ_{ij}^* denote the first principal angle between disjoint subspaces \mathcal{S}_i and \mathcal{S}_j , and let ϵ denote the maximal covering diameter of the points in \mathcal{Y}_i . A sufficient condition for EFS to occur for all points in \mathcal{Y}_i is that*

$$\max_{j \neq i} \cos(\theta_{ij}^*) < \frac{\sqrt{1 - \epsilon^2/4}}{1 + \epsilon/\sqrt[4]{12}}. \quad (15)$$

3.3 Connections to Previous Work

In this section, we will connect our results for OMP with previous analyses of EFS with BP provided in (Elhamifar and Vidal, 2010; Soltanolkotabi and Candès, 2012) for disjoint and intersecting subspaces respectively. Following this, we will contrast the geometry underlying EFS with exact recovery conditions used in standard applications of sparse recovery methods (Tropp, 2004, 2006).

3.3.1 SUBSPACE CLUSTERING WITH BP

In (Elhamifar and Vidal, 2010), the authors develop the following sufficient condition for EFS to occur for BP from a union of disjoint subspaces,

$$\max_{j \neq i} \cos(\theta_{ij}^*) < \max_{\tilde{Y}_i \in \mathbb{W}_i} \frac{\sigma_{\min}(\tilde{Y}_i)}{\sqrt{k_i}}, \quad (16)$$

where \mathbb{W}_i is the set of all full rank sub-matrices $\tilde{Y}_i \in \mathbb{R}^{n \times k_i}$ of the data matrix $Y_i \in \mathbb{R}^{n \times k_i}$ and $\sigma_{\min}(\tilde{Y}_i)$ is the minimum singular value of the sub-matrix \tilde{Y}_i . Since we assume that all of the data points have been normalized, $\sigma_{\min}(\tilde{Y}_i) \leq 1$; thus, the best case result that can be obtained is that the minimum principal angle, $\cos(\theta_{ij}^*) < 1/\sqrt{k_i}$. This suggests that the minimum principal angle of the union must go to zero, i.e., the union must consist of orthogonal subspaces, as the subspace dimension increases.

In contrast to the result in (16), our result for disjoint subspaces in Cor. 1 does not depend on the subspace dimension. Rather, our result requires that there are enough points in each subspace to achieve a sufficiently small covering; in which case, EFS can be guaranteed for disjoint subspaces of any dimension.

In (Soltanolkotabi and Candès, 2012), the authors develop the following sufficient condition for EFS with BP from unions of intersecting subspaces:

$$\mu_v(\mathcal{Y}_i) = \max_{y \in \mathcal{Y}_{-i}} \|V^{(i)T} y\|_\infty < r(\mathcal{Y}_i), \quad (17)$$

where the matrix $V^{(i)} \in \mathbb{R}^{d_i \times n}$ contains the dual directions (dual vector of each point in \mathcal{Y}_i embedded in \mathbb{R}^n) in its columns,³ and $r(\mathcal{Y}_i)$ is the inradius as defined in (10). In words, (17) requires that the maximum coherence between any point in \mathcal{Y}_{-i} and the dual directions contained in $V^{(i)}$ must be less than the inradius of the points in \mathcal{Y}_i .

Without going into the details of the interpretation of the dual directions $V^{(i)}$, the relationship between the coherence measure $\mu_v(\mathcal{Y}_i)$ provided in (Soltanolkotabi and Candès, 2012) and the mutual coherence $\mu_c(\mathcal{Y}_i)$ used in our analysis, is not apparent. However, when the points in each subspace cluster are distributed uniformly and at random along each subspace, the dual directions will also be distributed uniformly along each subspace.⁴ In this case, $\mu_v(\mathcal{Y}_i)$ will be roughly equivalent to the mutual coherence. When $\mu_c(\mathcal{Y}_i) \approx \mu_v(\mathcal{Y}_i)$, the result in (17) can be expressed in terms of the mutual coherence as $\mu_c(\mathcal{Y}_i) < r(\mathcal{Y}_i)$.

This simplification reveals the connection between the result in (Soltanolkotabi and Candès, 2012) for BP and the condition in Thm. 1 for OMP. In particular, our result for OMP requires that the mutual coherence is smaller than the inradius minus an additional term that depends on the product of the minimum principal angle between subspaces and a term that is linear in the covering diameter ϵ . Thus, the deterministic guarantee provided in (Soltanolkotabi and Candès, 2012) suggests that EFS may be guaranteed for BP for larger covering diameters than those required for OMP. However, when the covering diameter $\epsilon \rightarrow 0$ (as the subspace becomes more densely sampled), the guarantees for BP and OMP are roughly equivalent (assuming that $\mu_v \approx \mu_c$).

While our current results for OMP are slightly more restrictive than those obtained for BP, since EFS is only sufficient and not necessary for subspace recovery, we maintain that OMP offers a powerful alternative to BP for endogenous sparse recovery. This is due in part to the reduced computational complexity of OMP; in addition, we provide empirical evidence that demonstrates that OMP outperforms BP both in terms of EFS and with respect to the final clustering performance of each method on real data. Our empirical findings suggest that our worst-case results for OMP are quite pessimistic and that EFS occurs in practice for even sparser samplings of subspaces in the ensemble and/or higher degrees of overlap between subspaces (larger mutual coherence).

3.3.2 EXACT RECOVERY CONDITIONS FOR SPARSE RECOVERY

To provide further intuition about EFS in endogenous sparse recovery, we will compare the geometry underlying the EFS condition with the geometry of the exact recovery condition (ERC) for sparse signal recovery in (Tropp, 2004, 2006).

In standard analyses of sparse recovery algorithms, one assumes that a signal $y \in \mathbb{R}^n$ has been synthesized from a linear combination of atoms from a particular sub-dictionary $\Phi_\Lambda \in \mathbb{R}^{n \times k}$. To guarantee exact support recovery for y , we must ensure that we can uniquely recover an approximation of y that consists solely of atoms from Φ_Λ . To guarantee that a representation does not contain any atoms outside of the correct support set Λ , Tropp (2004) introduced a general *exact recovery condition* (ERC) that is sufficient for support recovery using both BP and OMP.

3. See Def. 2.2 in Soltanolkotabi and Candès (2012) for a formal definition of the dual directions and insight into the geometry underlying their guarantees for EFS via BP.

4. This approximation is based upon personal correspondence with M. Soltankotabi, one of the authors that developed the results for EFS with BP in Soltanolkotabi and Candès (2012).

Theorem 2 (Tropp, 2004) *For any signal supported over the sub-dictionary Φ_Λ , exact support recovery is guaranteed for both OMP and BP if*

$$\text{ERC}(\Lambda) = \max_{i \notin \Lambda} \|\Phi_\Lambda^\dagger \varphi_i\|_1 < 1. \quad (18)$$

A geometric interpretation of the ERC is that it provides a measure of how far a projected atom outside of the optimal set Λ lies from the antipodal convex hull of the atoms in Λ . When a projected atom lies outside of the antipodal convex hull formed by the set of points in the sub-dictionary Φ_Λ , then the ERC condition is violated and support recovery is not guaranteed. For this reason, the ERC requires that the maximum coherence between the columns of Φ is sufficiently low or that Φ is *incoherent*.

While the ERC condition requires a *global incoherence* property on all of the columns of Φ , the EFS condition requires a *local incoherence* property. In particular, the EFS condition requires that the projection of atoms in an incorrect subspace cluster \mathcal{Y}_{-i} onto \mathcal{S}_i must be incoherent with deep holes in \mathcal{Y}_i along \mathcal{S}_i . In addition, we need that the points within a subspace cluster are as coherent as possible in order to produce a small covering radius.

4. EFS for Uniformly Bounded Unions of Subspaces

In this section, we study the connection between EFS and the higher-order principal angles (beyond the minimum angle) between pairs of intersecting subspaces.

4.1 Subspace Distances

To characterize the “distance” between pairs of subspaces in the ensemble, the *principal angles* between subspaces will prove useful. As we saw in the previous section, the first principal angle θ_{ij}^* between subspaces \mathcal{S}_i and \mathcal{S}_j of dimension k_i and k_j is defined as the smallest angle between a pair of unit vectors (u_1, v_1) drawn from $\mathcal{S}_i \times \mathcal{S}_j$. The vector pair (u_1^*, v_1^*) that attains this minimum is referred to as the first set of principal vectors. The second principal angle is defined much like the first, except that the second set of principal vectors that define the second principal angle are required to be orthogonal to the first set of principal vectors (u_1^*, v_1^*) . The remaining principal angles are defined recursively in this way. The sequence of $k = \min(k_i, k_j)$ principal angles, $\theta_0 \leq \theta_1 \leq \dots \leq \theta_{k-1}$, is non-decreasing and all of the principal angles lie between $[0, \pi/2]$.

The definition above yields some insight into what the principal angles/vectors tell us about the geometry underlying a pair of subspaces; in practice, however, the principal angles are not computed in this recursive manner. Rather, a computationally efficient way to compute the principal angles between two subspaces is to first compute the singular values of the matrix $G = \Phi_i^T \Phi_j$, where $\Phi_i \in \mathbb{R}^{n \times k_i}$ is an ONB that spans subspace \mathcal{S}_i . We will refer to this set of singular values $\sigma_{ij} \in [0, 1]^k$ as the *cross-spectra* of the subspace pair $(\mathcal{S}_i, \mathcal{S}_j)$, where $k = \min(k_i, k_j)$ is the minimum dimension of the pair of subspaces of interest. The m^{th} smallest principal angle is related to the m^{th} largest value of the cross-spectra via the following relationship, $\cos(\theta_{ij}(m)) = \sigma_{ij}(m)$, where $\sigma_{ij}(m)$ is the m^{th} largest entry of the cross-spectra.

A pair of subspaces is said to be *disjoint* if the minimum principal angle is greater than zero. Non-disjoint or intersecting subspaces are defined as subspaces with minimum

principal angle equal to zero. The dimension of the intersection between two subspaces is equivalent to the number of entries of the cross-spectra that are equal to one. We also consider a slightly weaker notion of intersection that we call the *overlap* between two subspaces. The overlap q is defined as the rank(G) or equivalently, $q = \|\sigma_{ij}\|_0$. Thus, the overlap is greater than or equal to the dimension of the intersection between a pair of subspaces, $q \geq \dim(\mathcal{S}_i \cap \mathcal{S}_j)$.

4.2 Sufficient Conditions for EFS

The sufficient condition for EFS from disjoint subspaces in Cor. 1 reveals an interesting relationship between the covering radius and the minimum principal angle between pairs of subspaces in the ensemble. However, we have yet to reveal any dependence between EFS and higher-order principal angles. To make this connection more apparent, we will make additional assumptions about the distribution of points in the ensemble, namely that pairs of subspaces in the ensemble are *uniformly bounded* relative to the principal vectors supporting each pair of subspaces. The principal vectors can be interpreted as a pair of ONBs in \mathbb{R}^n that span \mathcal{S}_1 and \mathcal{S}_2 and are maximally coherent with one another.

To make this precise, consider the SVD of the matrix $G = \Phi_i^T \Phi_j = U \Sigma V^T$, where $\Phi_i \in \mathbb{R}^{n \times k_i}$ and $\Phi_j \in \mathbb{R}^{n \times k_j}$ are arbitrary ONBs that span subspaces \mathcal{S}_i and \mathcal{S}_j , respectively. Without loss of generality, assume that the dimension of \mathcal{S}_i is greater than or equal to the dimension of \mathcal{S}_j or that $k_i \geq k_j$. Let $\tilde{U} = \Phi_i U$ and $\tilde{V} = \Phi_j V$ denote the set of left and right principal vectors respectively, and $Y = [Y_i \ Y_j]$ be the matrix of data points, where Y_i and Y_j contain the points in \mathcal{S}_i and \mathcal{S}_j , respectively.

When the points in each subspace are incoherent with the principal vectors in the columns of \tilde{U} and \tilde{V} , we say that the ensemble Y is a *uniformly bounded union of subspaces*. Formally, we require the following incoherence property holds:

$$\left(\|Y_i^T \tilde{U}\|_\infty, \|Y_j^T \tilde{V}\|_\infty \right) \leq \gamma, \quad (19)$$

where $\|\cdot\|_\infty$ is the entry-wise maximum and $\gamma \in (0, 1]$. This property requires that the inner products between the principal vectors that span a particular subspace and the points in the subspace are all bounded by a fixed constant. In other words, for a union to have a small bounding constant, each point must have a dense (uniform) distribution with respect to the principal vectors that provide an ONB for the subspace.

When the points in each subspace are distributed such that (19) holds, we can rewrite the mutual coherence between any two points from different subspaces to reveal its dependence on higher-order principal angles. In particular, we show (in Section 7.2) that the coherence between the residual which is used to select the next point to be included in the representation, and a point in \mathcal{Y}_j at a particular iteration of the Alg. 1 is upper bounded by

$$\max_{y \in \mathcal{Y}_j} \frac{|\langle r, y \rangle|}{\|r\|_2} \leq \gamma \|\sigma_{ij}\|_1, \quad (20)$$

where γ is the bounding constant, σ_{ij} is the cross-spectra of the subspaces \mathcal{S}_i and \mathcal{S}_j , and $\|\sigma_{ij}\|_1$ is the ℓ_1 -norm of the cross-spectra or equivalently, the trace norm of $G = \Phi_i^T \Phi_j$. Using the bound in (20), we arrive at the following sufficient condition for EFS for uniformly bounded unions of subspaces.

Theorem 3 *Let Y be a uniformly bounded union of subspaces as defined in (19), where $r = \text{rank}(G)$, and $\gamma < \sqrt{1/r}$. Let σ_{ij} denote the cross-spectra of the subspaces \mathcal{S}_i and \mathcal{S}_j . A sufficient condition for EFS to occur for all of the points in \mathcal{Y}_i is that the covering diameter*

$$\epsilon < \min_{j \neq i} \sqrt{1 - \gamma^2 \|\sigma_{ij}\|_1^2}.$$

In words, this condition requires that the covering diameter of each subspace and the bounding constant of the union must both be sufficiently small in order to guarantee EFS. In order for an ensemble to have a small bounding constant, each point must have a balanced representation with respect to each the principal vectors that span their respective subspaces. To be precise, if we look at the coefficients β resulting from projecting each point $y \in \mathcal{S}_i$ onto the principal vectors than span \mathcal{S}_i , where $\beta = \tilde{U}^T y$, then to reduce the bounding constant, we need the energy in β to be spread as much as possible.

Our worst-case analysis assumes that the nonzero entries of the cross-spectra are equal, and thus each pair of supporting principal vectors are equally important in determining whether the ensemble will admit EFS. However, this assumption is not true in general. When the union is supported by principal vectors with non-uniform principal angles, our analysis suggests that a weaker form of incoherence is required. Instead of requiring incoherence with all principal vectors, the data must simply be incoherent with the principal vectors that correspond to small principal angles. This means that as long as points are not concentrated along the principal directions with small principal angles (i.e., intersections), then EFS can be guaranteed, even when subspaces exhibit non-trivial intersections. To test this prediction, we will study a *bounded energy model* for the data in Section 5.2. We show that when the dataset is sparsely sampled (larger covering radius), reducing the amount of energy that points contain in the intersections between two subspaces, increases the probability that points admit EFS dramatically.

Finally, our analysis of bounded unions suggests that the decay of the cross-spectra is likely to play an important role in determining whether points will admit EFS or not. Moreover, unions with different cross-spectral decay properties are likely to behave differently in terms of their respective probability of admitting EFS. To test this hypothesis, we will study the role that the structure of the cross-spectra plays in EFS in Section 5.3.

5. Experimental Results

In our theoretical analysis of EFS in Sections 3 and 4, we revealed an intimate connection between the covering radius of subspaces and the distribution of principal angles between pairs of subspaces in the ensemble. In this section, we will conduct an empirical study to explore these connections further. In particular, we will study the probability of EFS as we vary the covering radius as well as the dimension of the intersection and/or overlap between subspaces. In addition, we will study the role that the: (i) structure of the cross-spectra and (ii) amount of energy points have in the intersection between subspaces, have on EFS as predicted in our discussion of Thm. 3.

5.1 Generative Model for Synthetic Data

In order to study EFS for unions of subspaces with varied cross-spectra, we will generate data from unions of overlapping *block-sparse signals*.

5.1.1 CONSTRUCTING SUB-DICTIONARIES

We construct a pair of sub-dictionaries as follows: Take two subsets Ω_1 and Ω_2 of k atoms from a dictionary D of d n -dimensional vectors (atoms) $\{d_m\}_{m=1}^d$, where $d_m \in \mathbb{R}^n$ and $|\Omega_1| = |\Omega_2| = k$. Let $\Psi \in \mathbb{R}^{n \times k}$ denote the subset of atoms indexed by the first set $\Psi = D_{\Omega_1}$, and let $\Phi \in \mathbb{R}^{n \times k}$ denote the subset of atoms indexed by the second set $\Phi = D_{\Omega_2}$. Our goal is to select Ψ and Φ such that $G = \Psi^T \Phi$ is diagonal, i.e., $\langle \psi_i, \phi_j \rangle = 0$, if $i \neq j$, where ψ_i is the i^{th} element in Ψ and ϕ_j is the j^{th} element of Φ . In this case, the cross-spectra is defined as $\sigma = \text{diag}(G)$, where $\sigma \in [0, 1]^k$. For each union, we fix the “overlap” or the rank of $G = \Psi^T \Phi$ to a constant between zero (orthogonal) and k (maximal overlap), where the overlap $q \in [0, k]$ and $\delta = q/k$ is the overlap ratio between a pair of subspaces.

To generate a pair of k -dimensional subspaces with a q -dimensional overlap, we can pair the elements from Ψ and Φ such that the i^{th} entry of the cross-spectra equals

$$\sigma(i) = \begin{cases} |\langle \psi_i, \phi_i \rangle| & \text{if } 1 \leq i \leq q, \\ 0 & \text{if } i = q + 1 \leq i \leq k. \end{cases}$$

We can leverage the banded structure of shift-invariant dictionaries, e.g., dictionary matrices with localized Toeplitz structure, to generate subspaces with structured cross-spectra as follows.⁵ First, we fix a set of k incoherent (orthogonal) atoms from our shift-invariant dictionary, which we place in the columns of Ψ . Now, holding Ψ fixed, we set the i^{th} atom ϕ_i of the second sub-dictionary Φ to be a shifted version of the i^{th} atom ψ_i of the dictionary Ψ . To be precise, if we set $\psi_i = d_m$, where d_m is the m^{th} atom in our shift-invariant dictionary, then we will set $\phi_i = d_{m+\Delta}$ for a particular shift Δ . By varying the shift Δ , we can easily control the coherence between ψ_i and ϕ_i . In Fig. 4, we show an example of one such construction for $k = q = 5$. Since $\sigma \in (0, 1]^k$, the worst-case pair of subspaces with overlap equal to q is obtained when we pair q identical atoms with $k - q$ orthogonal atoms. In this case, the cross-spectra attains its maximum over its entire support and equals zero otherwise. For such unions, the overlap q equals the dimension of the intersection between the subspaces. We will refer to this class of block-sparse signals as *orthoblock sparse signals*.

5.1.2 COEFFICIENT SYNTHESIS

To synthesize a point that lives in the span of the sub-dictionary $\Psi \in \mathbb{R}^{n \times k}$, we combine the elements $\{\psi_1, \dots, \psi_k\}$ and subspace coefficients $\{\alpha(1), \dots, \alpha(k)\}$ linearly to form

$$y_i = \sum_{j=1}^k \psi_j \alpha(j),$$

5. While shift-invariant dictionaries appear in a wide range of applications of sparse recovery (Mailh  et al., 2008; Dyer et al., 2010), we introduce the idea of using shift-invariant dictionaries to create structured unions of subspaces for the first time here.




Figure 4: *Generating unions of subspaces from shift-invariant dictionaries.* An example of a collection of two sub-dictionaries of five atoms, where each of the atoms have a non-zero inner product with one other atom. This choice of sub-dictionaries produces a union of disjoint subspaces, where the overlap ratio $\delta = q/k = 1$.

where $\alpha(j)$ is the subspace coefficient associated with the j^{th} column in Ψ . Without loss of generality, we will assume that the elements in Ψ are sorted such that the values of the cross-spectra are monotonically decreasing. Let $y_i^c = \sum_{j=1}^q \psi_j \alpha_i(j)$ be the “common component” of y_i that lies in the intersection between (Φ, Ψ) and let $y_i^d = \sum_{j=q+1}^k \psi_j \alpha_i(j)$ denote the “disjoint component” of y_i that lies in the space orthogonal to the intersection between the subspaces.

For our experiments, we consider points drawn from one of the two following coefficient distributions.

1. *Uniformly Distributed on the Sphere:* Generate subspace coefficients according to a standard normal distribution and map the resulting point to the unit sphere

$$y_i = \frac{\sum_j \psi_j \alpha(j)}{\|\sum_j \psi_j \alpha(j)\|_2}, \quad \text{where } \alpha_i(j) \sim \mathcal{N}(0, 1).$$

2. *Bounded Energy Model:* Generate subspace coefficients uniformly on the sphere and rescale each coefficient in order to bound the energy in the common component

$$y_i = \frac{\tau y_i^c}{\|y_i^c\|_2} + \frac{(1 - \tau) y_i^d}{\|y_i^d\|_2}.$$

By restricting the total energy that each point has in its common component, the bounded energy model can be used to produce ensembles with small bounding constant to test the predictions in Thm. 3.

5.2 Phase Transitions for EFS

The goal of our first experiment is to study the probability of EFS—the probability that a point in the ensemble admits exact features—as we vary both the number and distribution of points in each subspace as well as the dimension of the intersection between subspaces. For this set of experiments, we generate a union of orthoblock sparse signals, where the overlap equals the dimension of the intersection.

Along the top row of Fig. 5, we display the probability of EFS for points drawn from a uniform distribution on the sphere: the probability of EFS is computed as we vary the *overlap ratio* $\delta = q/k \in [0, 1]$ in conjunction with the *oversampling ratio* $\rho = k/d \in [0, 1]$,




Figure 5: *Probability of EFS for different coefficient distributions.* The probability of EFS for a union of two subspaces of dimension $k = 20$ (left column) and $k = 50$ (right column). The probability of EFS is displayed as a function of the overlap ratio $\delta \in [0, 1)$ and the logarithm of the oversampling ratio $\log(\rho)$ (top row) and the mutual energy $\tau = \|y_c\|_2$ (bottom row) .

where $q = \text{rank}(\Phi_1^T \Phi_2)$ is equal to the dimension of the intersection between the subspaces, and d is the number of points per subspace. Along the bottom row of Fig. 5, we display the probability of EFS for points drawn from a bounded energy model: the probability of EFS is computed as we vary the overlap ratio δ and the mutual energy $\tau \in [0, 1)$. For these experiments, the subspace dimension is set to $k = 20$ (left) and $k = 50$ (right). To see the phase boundary that arises when we approach critical sampling (i.e., $\rho \approx 1$), we display our results in terms of the logarithm of the oversampling ratio. For these experiments, the results are averaged over 500 trials.

As our theory predicts, the oversampling ratio has a strong impact on the degree of overlap between subspaces that can be tolerated before EFS no longer occurs. In particular, as the number of points in each subspace increases (covering radius decreases), the probability of EFS obeys a second-order phase transition, i.e., there is a graceful degradation in the probability of EFS as the dimension of the intersection increases. When the pair of subspaces are densely sampled, the phase boundary is shifted all the way to $\delta = 0.7$. This is due to the fact that as each subspace is sampled more densely, the covering radius becomes sufficiently small to ensure that even when the overlap between planes is high, EFS still occurs with high probability. In contrast, when the subspaces are critically sampled, i.e., the number of points per subspace $d \approx k$, only a small amount of overlap can be toler-




Figure 6: *Probability of EFS for unions with structured cross-spectra.* Along the top row, we show the cross-spectra for different unions of block-sparse signals. Along the bottom row, we show the probability of EFS as we vary the overlap ratio $\delta \in [0, 1]$ for OMP (solid) and NN (dash).

ated, where $\delta < 0.1$. In addition to shifting the phase boundary, as the oversampling ratio increases, the width of the transition region increases.

Along the bottom row of Fig. 5, we study the impact of bounding the mutual energy of points in the ensemble as discussed in Section 4.2. In this experiment, we fix the oversampling ratio to $\rho = 0.1$ and vary the mutual energy τ in conjunction with the dimension of the intersection. By reducing the mutual energy of the union, the phase boundary for the uniformly distributed data in the top row is shifted from $\delta = 0.45$ to $\delta = 0.7$ for both $k = 20$ and $k = 50$. This result confirms our predictions in the discussion of Thm. 3 that by reducing the amount of energy that points have in their subspace intersections EFS will occur for higher degrees of overlap. Another interesting finding is that, once the mutual energy τ reaches a threshold, the phase boundary remains constant and reducing the energy further has no impact on the phase transitions for EFS.

5.3 Comparison of Feature Selection with OMP and NN

In this section, we compare the probability of EFS via OMP with the feature sets obtained via nearest neighbors (NN). First, we compare the performance of these methods for unions with varying cross-spectra. Second, we compare the phase transitions for unions of orthoblock sparse signals as we vary the overlap and number of points per subspace.

For our experiments, we generate pairs of subspaces with structured cross-spectra as described in Section 5.1.1. The cross-spectra arising from three different unions of block-sparse signals are displayed along the top row of Fig. 6. On the left, we show the cross-spectra for a union of overlapping orthoblock sparse signals with overlap ratio $\delta = 0.75$, where $q = 15$ and $k = 20$. The cross-spectra obtained from pairing shifted Lorentzian and

exponential atoms are displayed in the middle and right columns, respectively. In Fig. 6, we show the probability of EFS for OMP and NN for each of these three subspace unions as we vary the overlap q . To do this, we generate subspaces by setting their cross-spectra equal to the first q entries equal to the cross-spectra in Fig. 6 and setting the remaining $k - q$ entries of the cross-spectra equal to zero. Each subspace cluster is generated by sampling $d = 100$ points uniformly and at random along each normalized subspace cluster, as described in Section 5.1.2.

The gap between the probability of EFS for OMP and NN is markedly different for each of the three unions. In the first union of orthoblock sparse signals, the probability of EFS for OMP lies strictly above that obtained for the NN method, but the gap between the performance of both methods is relatively small. In the second union, both methods maintain a high probability of EFS, with OMP admitting nearly perfect feature sets even when the overlap ratio is maximal. In the third union, we observe that the gap between EFS for OMP and NN is most pronounced. In this case, the probability of EFS for NN sets decreases to 0.1, while OMP admits a very high probability of EFS, even when the overlap ratio is maximal.

This study provides a number of interesting insights into the role that higher-order principal angles between subspaces play in feature selection for both sparse recovery methods and NN. In particular, we observe that when data is distributed uniformly with respect to all of the principal directions between a pair of subspaces and the cross-spectra is sub-linear, then EFS may be guaranteed with high probability for all points in the set provided the sampling density is sufficiently high. This is in agreement with the discussion of EFS bounded unions in Section 4.2. Moreover, these results further support our claims that in order to truly understand and predict the behavior of endogenous sparse recovery from unions of subspaces, we require a description that relies on the entire cross-spectra.

In Fig. 7, we display the probability of EFS for OMP (left) and sets of NN (right) as we vary the overlap and the oversampling ratio. For this experiment, we consider unions of orthoblock sparse signals living on subspaces of dimension $k = 50$ and vary $\rho \in [0.2, 0.96]$ and $\delta \in [1/k, 1]$. An interesting result of this study is that there are regimes where the probability of EFS equals zero for NN but occurs for OMP with a non-trivial probability. In particular, we find that, when the sampling of each subspace is sparse, the gap between OMP and NN increases and OMP significantly outperforms NN in terms of their probability of EFS. Our study of EFS for structured cross-spectra suggests that the gap between EFS for NN and OMP will be even more pronounced for cross-spectra with superlinear decay.

5.4 Face Illumination Subspaces

In this section, we compare the performance of sparse recovery methods, i.e., BP and OMP, with NN selection for unions of *illumination subspaces* arising from a collection of images of faces under different lighting conditions. By fixing the camera center and position of the persons face and capturing multiple images under different lighting conditions, the resulting images can be well-approximated by a 5-dimensional subspace (Ramamoorthi, 2002).

In Fig. 2, we show three examples of the subspace affinity matrices obtained with NN, BP, and OMP for a collection of two different faces under 64 different illumination conditions from the Yale Database B (Georghiades et al., 2001), where each image has been subsampled




Figure 7: *Phase transitions for OMP and NN.* The probability of EFS for orthoblock sparse signals for OMP (a) and NN (b) feature sets as a function of the oversampling ratio $\rho = k/d$ and the overlap ratio $\delta = q/k$, where $k = 20$.

to 48×42 pixels, with $n = 2016$. In all of the examples, the data is sorted such that the images for each face are placed in a contiguous block.

To generate the NN affinity matrices in the left column, we compute the normalized inner products between all points in the dataset and then threshold each row to select the $k = 5$ nearest neighbors to each point. To generate the affinity matrices for both OMP and BP, we compute the sparse representations of each point in the dataset with Alg. 1 with $k = 5$ and by solving the BP denoising (BPDN) problem in (4) with $\kappa = 0.05$, respectively. The resulting coefficient vectors are placed into the rows of a matrix C and then the final subspace affinity W is computed by symmetrizing the coefficient matrix, $W = |C| + |C^T|$. To compute the BP affinity matrix, we threshold each coefficient vector c_i^* and select the top $k = 5$ coefficients (in absolute magnitude) and set the remaining coefficients to zero. We solved BPDN for different values of the noise tolerance parameter κ and included the estimate that yielded the best clustering performance. In addition, we computed the subspace affinity matrix using the full coefficient vectors (not thresholded) obtained via BPDN and with the least-squares estimate for top k coefficients from the thresholded BP coefficients, both of which produced worse clustering results.

In Table 1, we display the percentage of points that resulted in EFS and the final clustering results for all pairs of $\binom{38}{2}$ subspaces in the Yale B database, where we use the normalized graph cuts algorithm in (Shi and Malik, 2000) on the subspace affinity matrix for each method. Along the top row, we display the mean and median percentage of points that resulted in EFS for the full dataset (all 64 illumination conditions), half of the dataset (32 illumination conditions selected at random in each trial), and a quarter of the dataset (16 illumination conditions selected at random in each trial). While sparse recovery methods admit EFS rates that are comparable to NN on the full dataset, we find that OMP and BP obtain higher rates of EFS than NN for both the half and quarter datasets. These results are in agreement with our experiments on synthetic data that suggest that sparse recovery methods outperform NN when the amount of oversampling is low. In the middle row, we display the final clustering performance of each method with normalized graph cuts. In all

		Full-data			Half-data			Quarter-data		
		OMP	ℓ_1	NN	OMP	ℓ_1	NN	OMP	ℓ_1	NN
EFS (in %)	Mean	55.48	54.84	60.72	39.74	38.8	39.10	21.5	25.2	15.14
	Median	55.91	55.91	63.5	39.68	39.06	39.06	20.00	25.00	12.50
Clustering error (in %)	Mean	1.43	2.92	22.03	4.88	13.40	37.42	9.89	17.23	40.15
	Median	0.78	0.78	15.63	1.56	7.81	42.19	6.25	12.50	43.75
NCut cost (ratio)	Mean	8.43	8.29	6.48	6.40	5.55	4.41	4.49	3.83	3.02
	Median	7.91	7.43	5.80	6.02	5.09	4.03	4.10	3.49	2.81

Table 1: *Classification and EFS rates for illumination subspaces.* Shown are the aggregate results obtained over $\binom{38}{2}$ pairs of subspaces.

cases, sparse recovery methods outperform NN. A surprising result is that OMP provides significant improvements over BP for both the half and quartered datasets.

Since the theoretical guarantees we obtain for EFS with OMP are more pessimistic than the guarantees obtained for BP in (Soltanolkotabi and Candès, 2012), it is not obvious why OMP provides better clustering performance than BP. To investigate the cost of segmenting subspace affinity matrices obtained via the three methods, we compute the ratio of the absolute sum of the edges between points in the correct subspace (same face) to the absolute sum of the edges connecting points from different subspaces (different faces). We display this ratio along the bottom row of Table 1 (labeled Ncut cost). In all of our experiments, OMP has the highest ratio of in-cluster to out-of-cluster energies; this might be one reason why graph cuts performs better on subspace affinity matrices obtained via OMP in comparison to NN and BP.

6. Discussion

In this section, we provide insight into the implications of our results for different applications of sparse recovery and compressive sensing. Following this, we end with some open questions and lines for future research.

6.1 “Data Driven” Sparse Approximation

The standard paradigm in signal processing and approximation theory is to compute a compact representation of a signal in a fixed and pre-specified basis or dictionary. In most cases, the dictionaries used to form these representations are designed according to some mathematical desiderata. A more recent approach has been to learn a dictionary from a collection of data, such that the data admit a sparse representation with respect to the learned dictionary (Olshausen and Field, 1997; Aharon et al., 2006).

The applicability and utility of endogenous sparse recovery in subspace learning draws into question whether we can use endogenous sparse recovery for other tasks, including approximation and compression. The question that naturally arises is, “do we design a dictionary, learn a dictionary, or use the data as a dictionary?” Understanding the advantages and tradeoffs between each of these approaches is an interesting and open question.

6.2 Learning Block-Sparse Signal Models

Block-sparse signals and other structured sparse signals have received a great deal of attention over the past few years, especially in the context of compressive sensing from structured unions of subspaces (Lu and Do, 2008; Blumensath and Davies, 2009) and in model-based CS (Baraniuk et al., 2010). In all of these settings, the fact that a class or collection of signals admit structured support patterns is leveraged in order to obtain improved recovery of sparse signals in noise and in the presence of undersampling.

To exploit such structure in sparse signals—especially in situations where the structure of signals or blocks of active atoms may be changing across different instances in time, space, etc.—the underlying subspaces must be learned directly from the data. The methods that we have described for learning union of subspaces from ensembles of data can be utilized in the context of learning block sparse and other structured sparse signal models.

6.3 Beyond Coherence

While the maximum and cumulative coherence (Tropp, 2004) provide measures of the uniqueness of sub-dictionaries that are necessary to guarantee uniqueness and exact recovery, our current study suggests that examining the principal angles formed from pairs of sub-dictionaries could provide an even richer description of the geometric properties of a dictionary. Thus, a study of the principal angles formed by different subsets of atoms from a dictionary might provide new insights into sparse recovery with coherent dictionaries and compressive sensing from structured matrices. Finally, our empirical results in Section 5.3 suggest that there might exist an intrinsic difference between sparse recovery from dictionaries that exhibit sublinear versus superlinear decay in their principal angles or cross-spectra. It would be interesting to explore whether these two “classes” of dictionaries exhibit different phase transitions for sparse recovery.

6.4 Discriminative Dictionary Learning

While dictionary learning was originally proposed for learning dictionaries that admit sparse representations of a collection of signals (Olshausen and Field, 1997; Aharon et al., 2006), dictionary learning has recently been employed for classification. To use learned dictionaries for classification, a dictionary is learned for each class of training signals and then a sparse representation of a test signal is formed with respect to each of the learned dictionaries. The idea is that the test signal will admit a more compact representation with respect to the dictionary that was learned from the class of signals that the test signal belongs to.

Instead of learning these dictionaries independently of one another, *discriminative dictionary learning* (Mairal et al., 2008; Ramirez et al., 2010), aims to learn a collection of dictionaries $\{\Phi_1, \Phi_2, \dots, \Phi_p\}$ that are incoherent from one another. This is accomplished by minimizing either the spectral or Frobenius norm of the matrix product $\Phi_i^T \Phi_j$ between pairs of dictionaries. This same approach is utilized in (Mailh   et al., 2012) to learn sensing matrices for CS that are incoherent with a learned dictionary.

There are a number of interesting connections between discriminative dictionary learning and our current study of EFS from collections of unions of subspaces. In particular, our study provides new insights into the role that the principal angles between two dictionaries

tell us about our ability to separate classes of data based upon their sparse representations. Our study of EFS from unions with structured cross-spectra suggests that the decay of the cross-spectra between different data classes provides a powerful predictor of the performance of sparse recovery methods from data living on a union of low-dimensional subspaces. This suggests that, in discriminative dictionary learning, it might be more advantageous to reduce the ℓ_1 -norm of the entire cross-spectra rather than simply minimizing the maximum coherence and Frobenius norm between points in different subspaces as in (Mairal et al., 2008) and (Ramirez et al., 2010) respectively. To do this, each class of data must first be embedded within a subspace, a ONB is formed for each subspace, and then the ℓ_1 - norm of the cross-spectra is minimized. An interesting and relevant question is how one might impose such a constraint in discriminative dictionary learning methods.

6.5 Open Questions and Future Work

In this paper, we have developed new geometric conditions sparse recovery with OMP from data living on unions of subspaces. While these conditions provide new insight into the interplay between the geometry of pairs of subspaces and the sampling of each subspace, our worst-case analysis provides a very pessimistic outlook on EFS when compared with our empirical results. This gap between the theory and experimental results suggests that it might be possible to tighten our guarantees for EFS with OMP, either by incorporating a probabilistic or random sampling model for the generation of points on each subspace data or by improving upon the current analysis.

While the worst-case analysis we provide for OMP is more pessimistic than results provided for BP in Soltanolkotabi and Candès (2012), our empirical results suggest that OMP provides comparable performance to BP in terms of EFS. In some cases, OMP actually outperforms BP in terms of both EFS and their final classification performance on real data. Understanding what causes this gap in the performance of BP and OMP is a very interesting direction for future research.

Other areas for future research include studying: the average-case behavior of sparse recovery methods from unions of subspaces, the performance of sparse recovery methods on noisy unions of subspaces,⁶ the gap between sparse recovery methods and NN methods for feature selection, and the role that the decay of the cross-spectra plays in EFS for both sparse recovery and NN methods.

7. Proofs

7.1 Proof of Theorem 1

Our goal is to prove that, if (11) holds, then it is sufficient to guarantee that EFS occurs for every point in \mathcal{Y}_k when OMP is used for feature selection. We will prove this by induction.

Consider the greedy selection step in OMP (see Alg. 1) for a point y_i which belongs to the subspace cluster \mathcal{Y}_k . Recall that at the m^{th} step of OMP, the point that is maximally correlated with the signal residual will be selected to be included in the feature set Λ . The

6. In (Soltanolkotabi et al., 2013), they provide a robust variant of the SSC algorithm that can handle the case where points on multiple subspaces have been corrupted with Gaussian noise.

normalized residual at the m^{th} step is computed as

$$r^m = \frac{(I - P_\Lambda)y_i}{\|(I - P_\Lambda)y_i\|_2}, \quad (21)$$

where $P_\Lambda = Y_\Lambda Y_\Lambda^\dagger \in \mathbb{R}^{n \times n}$ is a projector onto the subspace spanned by the points in the current feature set Λ , where $|\Lambda| = m - 1$.

To guarantee that we select a point from \mathcal{S}_k , we require that the following greedy selection criterion holds:

$$\max_{v \in \mathcal{Y}_k} |\langle r^m, v \rangle| > \max_{v \notin \mathcal{Y}_k} |\langle r^m, v \rangle|. \quad (22)$$

We will prove that this selection criterion holds at each step of OMP by developing an upper bound on the RHS (the maximum inner product between the residual and a point outside of \mathcal{Y}_k) and a lower bound on the LHS (the minimum inner product between the residual and a point in \mathcal{Y}_k).

First, we will develop the upper bound on the RHS. In the first iteration, the residual is set to the signal of interest (y_i). In this case, we can bound the RHS by the mutual coherence $\mu_c = \max_{i \neq j} \mu_c(\mathcal{Y}_i, \mathcal{Y}_j)$ across all other sets

$$\max_{y_j \notin \mathcal{Y}_k} |\langle y_i, y_j \rangle| \leq \mu_c.$$

Now assume that at the m^{th} iteration we have selected points from the correct subspace cluster. This implies that our signal residual still lies within the span of \mathcal{Y}_k , and thus we can write the residual $r^m = z + e$, where z is the closest point to r^m in \mathcal{Y}_k and e is the remaining portion of the residual which also lies in \mathcal{S}_k . Thus, we can bound the RHS as follows

$$\begin{aligned} \max_{y_j \notin \mathcal{Y}_k} |\langle r^m, y_j \rangle| &= \max_{y_j \notin \mathcal{Y}_k} |\langle z + e, y_j \rangle| \\ &\leq \max_{y_j \notin \mathcal{Y}_k} |\langle z, y_j \rangle| + |\langle e, y_j \rangle| \\ &\leq \mu_c + \max_{y_j \notin \mathcal{Y}_k} |\langle e, y_j \rangle| \\ &\leq \mu_c + \cos(\theta_0) \|e\|_2 \|y_i\|_2. \end{aligned}$$

Using the fact that $\text{cover}(\mathcal{Y}_k) = \epsilon/2$, we can bound the ℓ_2 -norm of the vector e as

$$\begin{aligned} \|e\|_2 &= \|r - z\|_2 \\ &= \sqrt{\|r\|_2^2 + \|z\|_2^2 - 2|\langle r, z \rangle|} \\ &\leq \sqrt{2 - 2\sqrt{1 - (\epsilon/2)^2}} \\ &= \sqrt{2 - \sqrt{4 - \epsilon^2}}. \end{aligned}$$

Plugging this quantity into our expression for the RHS, we arrive at the following upper bound

$$\max_{y_j \notin \mathcal{Y}_k} |\langle r^m, y_j \rangle| \leq \mu_c + \cos(\theta_0) \sqrt{2 - \sqrt{4 - \epsilon^2}} < \mu_c + \cos(\theta_0) \frac{\epsilon}{\sqrt[4]{12}},$$

where the final simplification comes from invoking the following Lemma.

Lemma 1 For $0 \leq x \leq 1$,

$$\sqrt{2 - \sqrt{4 - x^2}} \leq \frac{x}{\sqrt[4]{12}}.$$

Proof of Lemma 1: We wish to develop an upper bound on the function

$$f(x) = 2 - \sqrt{4 - x^2}, \quad \text{for } 0 \leq x \leq 1.$$

Thus our goal is to identify a function $g(x)$, where $f'(x) \leq g'(x)$ for $0 \leq x \leq 1$, and $g(0) = f(0)$. The derivative of $f(x)$ can be upper bounded easily as follows

$$f'(x) = \frac{x}{\sqrt{4 - x^2}} \leq \frac{x}{\sqrt{3}}, \quad \text{for } 0 \leq x \leq 1.$$

Thus, $g'(x) = x/\sqrt{3}$, and $g(x) = x^2/\sqrt{12}$; this ensures that $f'(x) \leq g'(x)$ for $0 \leq x \leq 1$, and $g(0) = f(0)$. By the Fundamental Theorem of Integral Calculus, $g(x)$ provides an upper bound for $f(x)$ over the domain of interest where, $0 \leq x \leq 1$. To obtain the final result, take the square root of both sides, $\sqrt{2 - \sqrt{4 - x^2}} \leq \sqrt{x^2/\sqrt{12}} = x/\sqrt[4]{12}$. \square

Second, we will develop the lower bound on the LHS of the greedy selection criterion. To ensure that we select a point from \mathcal{Y}_k at the first iteration, we require that y_i 's nearest neighbor belongs to the same subspace cluster. Let y_{nn}^i denote the nearest neighbor to y_i

$$y_{nn}^i = \arg \max_{j \neq i} |\langle y_i, y_j \rangle|.$$

If y_{nn}^i and y_i both lie in \mathcal{Y}_k , then the first point selected via OMP will result in EFS.

Let us assume that the points in \mathcal{Y}_k admit an ϵ -covering of the subspace cluster \mathcal{S}_k , or that $\text{cover}(\mathcal{Y}_k) = \epsilon/2$. In this case, we have the following bound in effect

$$\max_{y_j \in \mathcal{Y}_k} |\langle r^m, y_j \rangle| \geq \sqrt{1 - \frac{\epsilon^2}{4}}.$$

Putting our upper and lower bound together and rearranging terms, we arrive at our final condition on the mutual coherence

$$\mu_c < \sqrt{1 - \frac{\epsilon^2}{4}} - \cos(\theta_0) \frac{\epsilon}{\sqrt[4]{12}}.$$

Since we have shown that this condition is sufficient to guarantee EFS at each step of Alg. 1 provided the residual stays in the correct subspace, Thm. 1 follows by induction. \square

7.2 Proof of Theorem 3

To prove Thm. 3, we will assume that the union of subspaces is uniformly bounded in accordance with (19). This assumption enables us to develop a tighter upper bound on the mutual coherence between the residual $r \in \mathcal{S}_i$ and the points in \mathcal{Y}_j . Since $r \in \mathcal{S}_i$, the

residual can be expressed as $r = \Phi_i \alpha$, where Φ_i is an ONB that spans \mathcal{S}_i and $\alpha = \Phi_i^T r$. Similarly, we can write all of the points in \mathcal{Y}_j as $y = \Phi_j \beta$, where Φ_j is an ONB that spans \mathcal{S}_j and $\|\beta\|_2 = \|y\|_2 = 1$ because Φ_j is a unitary matrix that preserves the ℓ_2 -norm of y . Let $\beta_i = \Phi_j^T y_i$ denote the subspace coefficients obtained by expanding y_i in terms of y_j and let $\mathcal{B}_j = \{\beta_i\}_{i=1}^{d_j}$ denote the set of all subspace coefficients for all $y_i \in \mathcal{Y}_j$. The coherence between the residual and a point in a different subspace can be expanded as follows:

$$\begin{aligned}
 \max_{y \in \mathcal{Y}_j} \frac{|\langle r, y \rangle|}{\|r\|_2} &= \max_{\beta \in \mathcal{B}_j} \frac{|\langle \Phi_i \alpha, \Phi_j \beta \rangle|}{\|\alpha\|_2} \\
 &= \max_{\beta \in \mathcal{B}_j} \frac{|\langle \alpha, \Phi_i^T \Phi_j \beta \rangle|}{\|\alpha\|_2} \\
 &= \max_{\beta \in \mathcal{B}_j} \frac{|\langle \alpha, U \Sigma V^T \beta \rangle|}{\|\alpha\|_2} \\
 &= \max_{\beta \in \mathcal{B}_j} \frac{|\langle U^T \alpha, \Sigma V^T \beta \rangle|}{\|\alpha\|_2} \\
 &\leq \max_{\beta \in \mathcal{B}_j} \frac{\|U^T \alpha\|_\infty}{\|\alpha\|_2} \|\Sigma V^T \beta\|_1,
 \end{aligned} \tag{23}$$

where the last step comes from an application of Holder's inequality, i.e., $|\langle w, z \rangle| < \|w\|_\infty \|z\|_1$.

Now, we tackle the final expression in (23), which we can write as

$$\max_{\beta \in \mathcal{B}_j} \|\Sigma V^T \beta\|_1 = \max_{y \in \mathcal{Y}_j} \|\Sigma V^T \Phi_j^T y\|_1 = \max_{y \in \mathcal{Y}_j} \|\Sigma (\Phi_j V)^T y\|_1, \tag{24}$$

where the matrix $\Phi_j V$ contains the principal vectors in subspace \mathcal{S}_j . Since we have assumed that the union is bounded, this implies that the inner product between all principal vectors and the points in \mathcal{Y}_j are bounded by γ . It follows that $\|(\Phi_j V)^T y\|_\infty \leq \gamma$. Now, suppose that $\gamma^2 q < 1$, where $q = \text{rank}(G) = \|\sigma_{ij}\|_0$. In this case,

$$\max_{y \in \mathcal{Y}_j} \|\Sigma (\Phi_j V)^T y\|_1 \leq \gamma \|\sigma_{ij}\|_1. \tag{25}$$

Note that for bounded unions of subspaces, the term on the right can be made small by requiring that the bounding constant $\gamma \ll 1$. Plugging this bound into (23), we obtain the following expression

$$\max_{y \in \mathcal{Y}_j} \frac{|\langle r, y \rangle|}{\|r\|_2} \leq \gamma \|\sigma_{ij}\|_1 \frac{\|U^T \alpha\|_\infty}{\|\alpha\|_2} = \gamma \|\sigma_{ij}\|_1 \|U\|_{2,2} = \gamma \|\sigma_{ij}\|_1,$$

where this last simplification comes from the fact that U is unitary and has spectral norm equal to one. Note that this bound on the mutual coherence is informative only when $\gamma \|\sigma_{ij}\|_1 < \sigma_{\max}$. This completes the proof. \square

Acknowledgements

Thanks to Dr. Christoph Studer, Dr. Chinmay Hegde, and Mahdi Soltanolkotabi for helpful discussions and comments on this paper. Thanks also to Dr. Arian Maleki and Dr. Joel Tropp for helpful discussions. ED was supported by a NSF GRFP 0940902 and a Texas Instruments Distinguished Graduate Fellowship. ACS and RGB were partially supported by following grants: NSF CCF-1117939, CCF-0431150, CCF-0728867, CCF-0926127; DARPA N66001-11-1-4090, N66001-11-C-4092; ONR N00014-08-1-1112, N00014-10-1-0989; AFOSR FA9550-09-1-0432; ARO MURIs W911NF-07-1-0185 and W911NF-09-1-0383.

References

- M. Aharon, M. Elad, and A. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Trans. Signal Processing*, 54(11):4311–4322, 2006.
- E. Arias-Castro, G. Chen, and G. Lerman. Spectral clustering based on local linear approximations. *arXiv:1001.1323v2 [stat.ML]*, September 2010.
- R. G. Baraniuk, V. Cevher, M. Duarte, and C. Hegde. Model-based compressive sensing. *IEEE Trans. Inform. Theory*, 56(4):1982–2001, 2010.
- R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Anal. Machine Intell.*, 25(2):218–233, February 2003.
- T. Blumensath and M. Davies. Sampling theorems for signals from the union of finite-dimensional linear subspaces. *IEEE Trans. Inform. Theory*, 55(4):1872–1882, 2009.
- G. Chen and G. Lerman. Spectral curvature clustering. *Int. J. Computer Vision*, 81:317–330, 2009.
- S. Chen, D. Donoho, and M. Saunders. Atomic decomposition by basis pursuit. *SIAM J. Sci. Comp.*, 20(1):33–61, 1998.
- G. Davis, S. Mallat, and Z. Zhang. Adaptive time-frequency decompositions. *SPIE J. Opt. Engin.*, 33(7):2183–2191, 1994.
- E. L. Dyer. Endogenous sparse recovery. Master’s thesis, Electrical & Computer Eng. Dept., Rice University, 2011.
- E. L. Dyer, M. Duarte, D. J. Johnson, and R. G. Baraniuk. Recovering spikes from noisy neuronal calcium signals via structured sparse approximation. *Proc. Int. Conf. on Latent Variable Analysis and Sig. Separation*, pages 604–611, 2010.
- E. Elhamifar and R. Vidal. Sparse subspace clustering. In *Proc. IEEE Conf. Comp. Vis. Patt. Recog. (CVPR)*, June 2009.
- E. Elhamifar and R. Vidal. Clustering disjoint subspaces via sparse representation. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, pages 1926–1929, March 2010.

- E. Elhamifar and R. Vidal. Sparse subspace clustering: algorithm, theory, and applications. *IEEE Trans. Pattern Anal. Machine Intell.*, 2012.
- A. S. Georghiades, P. N. Bellhumeur, and D. J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Machine Intell.*, 23(6):643–660, 2001.
- B. V. Gowreesunker, A. Tewfik, V. Tadipatri, J. Ashe, G. Pellize, and R. Gupta. A subspace approach to learning recurrent features from brain activity. *IEEE Trans. Neur. Sys. Reh.*, 19(3):240–248, 2011.
- K. Kanatani. Motion segmentation by subspace separation and model selection. In *Proc. IEEE Int. Conf. Comp. Vis. (ICCV)*, 2001.
- Y. Lu and M. Do. Sampling signals from a union of subspaces. *IEEE Sig. Proc. Mag.*, 25(2):41–47, March 2008.
- B. Mailh , S. Lesage, R. Gribonval, F. Bimbot, and P. Vandergheynst. Shift-invariant dictionary learning for sparse representations: extending K-SVD. In *Proc. Europ. Sig. Processing Conf. (EUSIPCO)*, 2008.
- B. Mailh , D. Barchiesi, and M. D. Plumbley. INK-SVD: Learning incoherent dictionaries for sparse representations. In *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Processing (ICASSP)*, pages 3573–3576, march 2012.
- J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Discriminative learned dictionaries for local image analysis. In *Proc. IEEE Conf. Comp. Vis. Patt. Recog. (CVPR)*, June 2008.
- B. Olshausen and D. Field. Sparse coding with an overcomplete basis set: a strategy employed by V1. *Vision Res.*, 37:3311–3325, 1997.
- R. Ramamoorthi. Analytic PCA construction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE Trans. Pattern Anal. Machine Intell.*, 24(10):1322–1333, 2002.
- I. Ramirez, P. Sprechmann, and G. Sapiro. Classification and clustering via dictionary learning with structured incoherence and shared features. In *Proc. IEEE Conf. Comp. Vis. Patt. Recog. (CVPR)*, pages 3501–3508, June 2010.
- J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.*, 22(8):888–905, August 2000.
- M. Soltanolkotabi and E. J. Cand s. A geometric analysis of subspace clustering with outliers. *Annals of Statistics*, 40(4):2195–2238, 2012.
- M. Soltanolkotabi, E. Elhamifar, and E. J. Cand s. Robust subspace clustering. *CoRR*, abs/1301.2603, 2013.
- J. A. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Trans. Inform. Theory*, 50(10):2231–2242, 2004.

- J. A. Tropp. Just relax: convex programming methods for identifying sparse signals in noise. *IEEE Trans. Inform. Theory*, 52(3):1030–1051, March 2006.
- J. Yan and M. Pollefeys. A general framework for motion segmentation: Independent, articulated, rigid, non-rigid, degenerate and non-degenerate. In *Proc. European Conf. Comp. Vision (ECCV)*, 2006.
- T. Zhang, A. Szlam, Y. Wang, and G. Lerman. Hybrid linear modeling via local best-fit flats. *arXiv:1010.3460v1 [cs.CV]*, October 2010.